

CLAIMS

What is claimed is:

1 1. An input/output controller integrated circuit
2 comprising:
3 a host interface subsystem to couple to a host to
4 receive host commands and to transceive data blocks
5 with the host in response to the host commands;

6 and

7 a mapping controller coupled to the host interface
8 subsystem, the mapping controller to map logical block
9 addresses of the host commands into physical block
10 addresses of one or more peripherals to transceive the
11 data blocks with the one or more peripherals.

1 2. The integrated input/output controller of claim 1

2 further comprising:

3 a peripheral interface subsystem to couple to the one or
4 more peripherals to issue peripheral commands and to
5 transceive data blocks with the one or more peripherals in
6 response to the one or more peripheral commands.

1 3. An integrated input/output (I/O) controller integrated
2 into an integrated circuit to read and write data between a host
3 and one or more peripherals, the integrated I/O controller
4 including circuits to perform:

5 receiving a high level I/O command from a host;
6 parsing the high level I/O command to determine whether
7 to read or write data;
8 mapping the high level I/O request into one or more
9 peripheral I/O commands, the one or more peripheral I/O
10 commands indicating which of the one or more peripherals and
11 which respective data locations are to be accessed; and
12 servicing the high level I/O request by reading or
13 writing data between the host and the one or more peripherals
14 using the respective data locations.

1 4. The integrated I/O controller of claim 3, further
2 including circuits to perform:
3 prior to servicing the high level I/O request, storing
4 data temporarily into an external cache buffer from the data
5 flow between the host and the one or more peripherals.

1 5. The integrated I/O controller of claim 3, wherein,
2 each of the one or more peripheral I/O commands further
3 indicates the number of blocks of data to be serviced.

1 6. The integrated I/O controller of claim 3, wherein,
2 mapping the high level I/O request into one or more
3 peripheral I/O commands includes circuits to perform
4 parsing a high level command from an I/O request
5 packet,

decoding the high level command and generating a range operation request in response thereto, and generating the one or more peripheral I/O commands in response to the range operation request.

7. The integrated I/O controller of claim 3, wherein,
each of the one or more peripherals are storage disks
and each of the one or more peripheral I/O commands are Small
Computer System Interface (SCSI) disk I/O commands.
8. The integrated I/O controller of claim 3, wherein,
the high level I/O command is a command of a Small
Computer System Interface (SCSI) Command Descriptor Block
(CDB) standard.
9. The integrated I/O controller of claim 3, wherein,
the circuits of the integrated I/O controller are hard
wired circuits.
10. The integrated I/O controller of claim 3, wherein,
the circuits of the integrated I/O controller are
microcoded circuits and state machines operating
concurrently.
11. The integrated I/O controller of claim 3, wherein,

2 the circuits of the integrated I/O controller are hard
3 wired circuits, microcoded circuits and state machines
4 operating concurrently.

1 12. The integrated I/O controller of claim 3, wherein,
2 the circuits of the integrated I/O controller are
3 programmable micro-controllers operating concurrently.
QH

1 13. An input/output controller integrated into a single
2 integrated circuit comprising:

3 a host interface subsystem to couple to a host to
4 receive host commands and to transceive data blocks
5 with the host in response to the host commands;

6 a peripheral interface subsystem to couple to a
7 peripheral to issue peripheral commands and to
8 transceive data blocks with the peripheral in response
9 to the peripheral commands; and

10 a mapping controller coupled to the host interface
11 subsystem and the peripheral interface subsystem, the
12 mapping controller to map to blocks of data storage of
13 the peripheral.

1 14. The integrated input/output controller of claim 9

2 further comprising:

3 a micro-controller subsystem coupled to the host
4 interface subsystem and the peripheral interface subsystem,

5 the micro-controller subsystem to perform initialization and
6 to process errors and exception events.

15. The integrated input/output controller of claim 13

? further comprising:

er comprising:
a cache manager coupled to the host interface subsystem
and the peripheral interface subsystem, the cache manager to
manage entries in a cache buffer to temporarily store data of
the data flow between the peripheral and the host.

16. The integrated input/output controller of claim 15

further comprising:

a buffer manager coupled to the host interface subsystem, the peripheral interface subsystem, and the cache manager, the buffer manager to manage data storage in the cache buffer.

17. The integrated input/output controller of claim 13

2 wherein,

2 wherein,
3 the host interface subsystem includes a fibre channel
4 host port to transceive data with the host using a fibre
5 channel protocol.

18. The integrated input/output controller of claim 13

2 wherein,

the host interface subsystem includes

a host exchange controller to control the physical connection and protocol of the host.

19. The integrated input/output controller of claim 13

wherein, includes

the host interface subsystem includes a command decode controller to decode host commands.

20. The integrated input/output controller of claim 13

wherein,
the peripheral interface subsystem includes
a peripheral exchange controller to control the
physical connection and protocol of the peripheral.

21. The integrated input/output controller of claim 13

wherein,
the peripheral interface subsystem includes a
Fibrechannel disk port to transceive data with the peripheral
using a Fibrechannel protocol.

1 22. A method of processing host write commands between a
2 host and one or more peripherals, the method comprising:
3 receiving a high level write command from a host, the
4 high level write command requesting to write data from the
5 host transparently to one or more peripherals;

6 mapping the high level write command into one or more
7 low level peripheral write commands, the one or more low
8 level peripheral write commands indicating data storage
9 locations in the one or more peripherals into which data from
10 the host is to be written;

11 servicing the high level write command by receiving data
12 from the host for storage into the data storage locations in
13 the one or more peripherals; and

14 writing the data from the host into the data storage
15 locations in the one or more peripherals in response to the
16 one or more low level peripheral write commands.

17 23. The method of claim 22 further comprising:
18 prior to writing the data from the host into the data
19 storage locations, temporarily storing the data from the host
20 into a buffer memory.

21 24. The method of claim 22 wherein,
22 the high level write command indicates the number of
23 blocks of data to be written from the host.

24 25. A method of processing host read commands between a host
25 and one or more peripherals, the method comprising:
1 receiving a high level read command from a host, the
2 high level read command requesting to transparently read data
3 from the one or more peripherals to the host;
4

6 mapping the high level read command into one or more low
7 level peripheral read commands, the one or more low level
8 peripheral read commands indicating data storage locations in
9 the one or more peripherals from which data is to be accessed
10 and read out to the host;

11 servicing the high level read request by reading data
12 from the data storage locations in the one or more
13 peripherals in response to the one or more low level
14 peripheral read commands; and

15 providing the data to the host.

16 26. The method of claim 25 further comprising:
17 prior to providing the data to the host, temporarily
18 storing the data from the one or more peripherals into a
19 buffer memory.

20 27. The method of claim 25 wherein,
21 the high level read command indicates the number of
22 blocks of data to be read into the host.

23 28. An integrated input/output controller comprising:
24 a semiconductor integrated circuit including,
25 a host interface to couple to one or more
26 servers to transceive data between the one or more
27 servers and the integrated input/output controller in
28 response to host commands,

a mapping controller coupled to the host interface subsystem, the mapping controller to map to blocks of data storage of a peripheral, and a micro-controller coupled to the host interface subsystem, the micro-controller to perform initialization and handle error and exception handling events.

29. The integrated input/output controller of claim 28

wherein,
the host interface subsystem includes a fibre channel host port to transceive data with the one or more servers using a fibre channel protocol.

30. The integrated input/output controller of claim 28

wherein,
the host interface subsystem includes a command decode controller to receive and decode host command packets.

1 31. The integrated input/output controller of claim 30

2 wherein,
3 the command decode controller to further validate a host
4 command and to initiate execution of the host command by the
5 integrated input/output controller.

32. The integrated input/output controller of claim 30

2 wherein,

3 the command decode controller to further validate a host
4 command and to queue the host command in a queue to be
5 executed in an order, the command queue associated with a
6 volume accessible by one of the one or more servers.
BS

1 33. The integrated input/output controller of claim 30
2 wherein,

3 the command decode controller to further validate a host
4 command and to determine that the host command is invalid,
5 the host command is passed to the micro-controller subsystem
6 to process the invalidity.

1 34. The integrated input/output controller of claim 30
2 wherein,

3 the host command packets are high level input/output
4 requests.

1 35. The integrated input/output controller of claim 28
2 further comprising:

3 a peripheral interface to couple to one or more
4 peripheral devices and transceive data between the one or
5 more peripheral devices and the one or more servers coupled
6 to the host port.

1 36. The integrated input/output controller of claim 28
2 wherein,

6 controller further comprising
7 a disk interface to couple to the one or more
8 drives.

[Signature] a disk interface
7 disks of at least one disk array to transceive the
8 data to or from the one or more servers; and
9

the mapping controller is a RAID mapping controller to flexibly control the mapping of blocks of data storage on the one or more disks of the at least one disk array.

37. The integrated input/output controller of claim 36

wherein,
the one or more disks of the at least one array of disks
are magnetic storage media, optical storage media or
semiconductor storage media.

38. The integrated input/output controller of claim 36

2 wherein,
3 the disk interface subsystem includes one or more fibre
4 channel disk ports to transceive data with the one or more
5 disks of the at least one disk array using a fibre channel
6 protocol.

1 39. The integrated input/output controller of claim 28

? wherein,

005127.P001
Express Mail No.: EL466330618US

The mapping controller provides RAID mapping automation.

41. The integrated input/output controller of claim 28
wherein,
the mapping controller to receive a requested command input packet to generate expanded command output packets in response thereto, the requested command input packet functions as a logical address and the expanded command output packets function as physical addresses.

1 42. The integrated input/output controller of claim 28
2 further comprising:
3 a buffer manager and a cache manager, the buffer m
4 and the cache manager to couple to a cache buffer and a
5 table buffer respectively to flexibly control the ready
6 writing of data to and from the mapped data storage on
7 one or more disks of the at least one disk array.

1 43. A storage area network for central data storage and
2 management, the storage area network comprising:
3 at least one server to couple to a network;

4 at least one disk array having a plurality of disks; and
5 at least one fibre channel controller to couple to at
6 least one server and at least one disk array, the at least
7 one fibre channel controller to read and write data between
8 the at least one server and the at least one disk array, the
9 at least one fiber channel controller having
10 a cache memory,
11 a microprocessor to initialize the fiber channel
12 controller upon power up and reset,
13 a programmable random access memory (PRAM) to store
14 initialization instructions in firmware, and
15 a hardware redundant array of independent disks (RAID)
16 controller to control the reading and writing of data between
17 the at least one server and the at least one disk array.

1 44. The storage area network of claim 43 for central data
2 storage and management, wherein the hardware RAID controller
3 includes,

4 a host interface to couple to the at least one
5 server to receive data from the at least one server for
6 storage into the one or more disks of the at least one
7 disk array and to transmit data to the at least one
8 server for data accessed from the one or more disks of
9 the at least one disk array;
10 a disk interface to couple to the one or more
11 disks of the at least one disk array to transmit data
12 to the one or more disks of the at least one disk array

13 for storage and to receive data from the one or more
14 disks of the at least one disk array when accessed,
15 a RAID mapping controller to flexibly control the
16 mapping of blocks of storage on the one or more disks
17 of the at least one disk array, and
18 a micro-controller coupled to the host interface
19 subsystem, the disk interface subsystem, and the RAID
20 mapping subsystem, the micro-controller to handle non-
21 data flow commands, error and exception handling events
22 as well as system initialization.

BB

44. The storage area network of claim 43 for central data
storage and management, wherein,

the cache memory to provide a cache data buffer and a
cache table buffer for the integrated RAID controller
integrated circuit.

46. The storage area network of claim 43 for central data
storage and management, wherein,

the plurality of disks of the at least one array of
disks are magnetic storage media, optical storage media or
semiconductor storage media.

47. An integrated input/output (I/O) controller comprising:
a semiconductor integrated circuit including,
a host command manager to manage high level host
I/O requests from a plurality of hosts;

a mapping engine to map high level host I/O requests into low level I/O commands; and a low level command manager to manage data read and data write accesses into and out of a peripheral device in response to the low level I/O commands.

48. The integrated input/output (I/O) controller of claim 47

further comprising:

prising:
a buffer manager to arbitrate access by the one or
more servers and to control data reads and data writes
into and out of a buffer memory.

49. The integrated input/output (I/O) controller of claim 47

further comprising:

a micro-controller to handle non-data flow commands, system initialization and error handling exception conditions.

50. The integrated input/output (I/O) controller of claim 47

2 wherein

the integrated I/O controller is a RAID controller
and the peripheral device is a plurality of disks
responsive to disk I/O commands.

3 receiving an I/O request from at least one server, the
4 I/O request indicating a volume to access;

5 mapping the I/O request into a disk I/O command, the
6 disk I/O command indicating one or more disks of the
7 plurality of disks to process the disk I/O command; and

8 managing a data read or a data write access to the
9 plurality of disks in response to the disk I/O command.

1 52. The method of claim 51 for transceiving data to a
2 plurality of disks in an integrated circuit further comprising:

3 arbitrating a data read or a data write access to a
4 buffer memory to temporarily store the data prior to the data
5 read or data write access to the plurality of disks.